





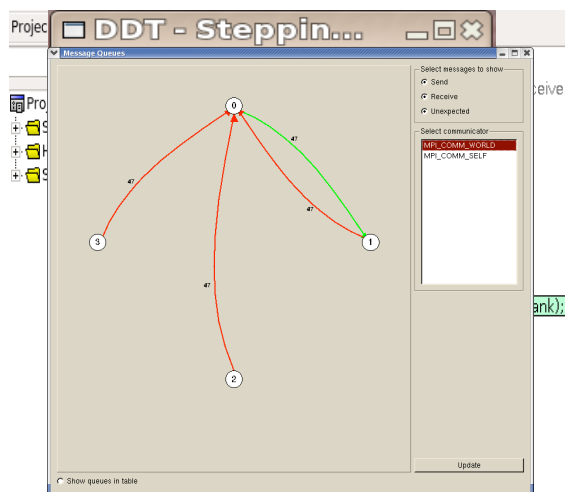
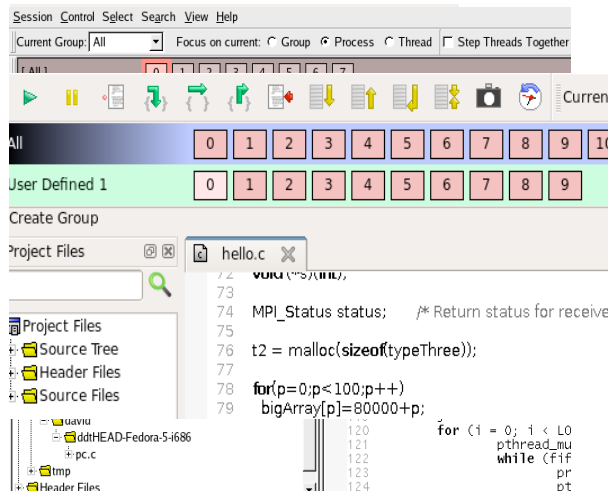
Debugging the Future with DDT at ORNL

David Lecomber
david@allinea.com
CTO

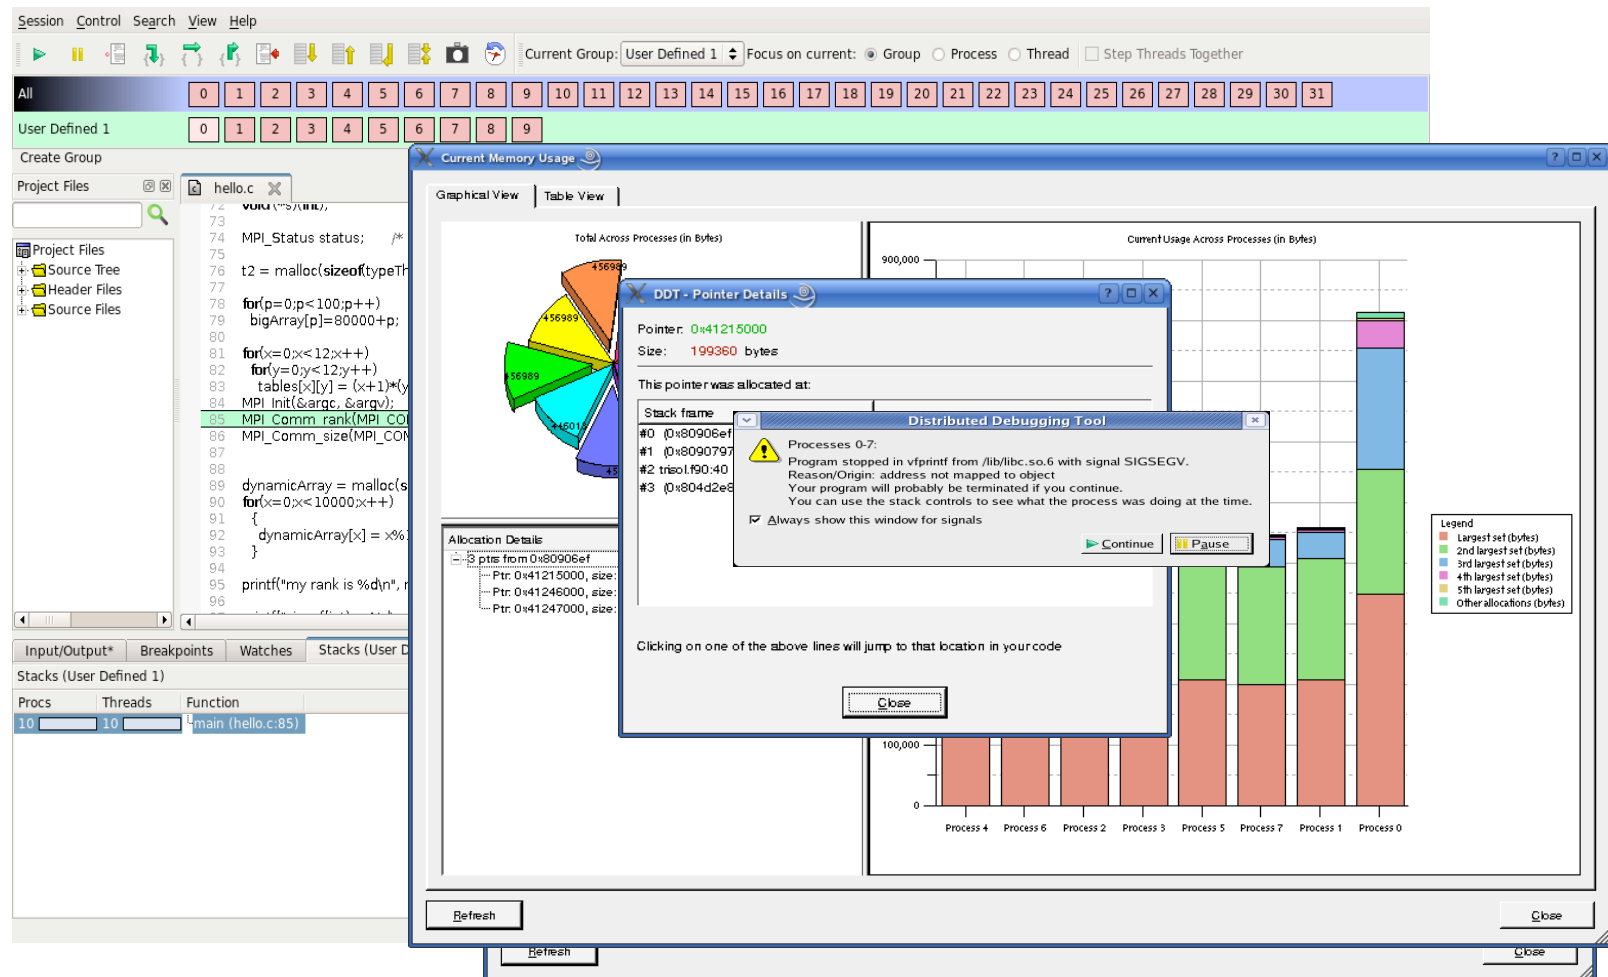
- Processor counts growing rapidly
- GPUs entering HPC
- Large hybrid systems imminent
- But what happens when software doesn't work?



- A company focussed on HPC tools
 - DDT – the easiest tool for debugging parallel codes at every scale
 - OPT – instruments code to find bottlenecks
 - DDTLite – plug-in for Visual Studio 2008
- Significance of previous graph?
 - **Everyone** aspires to have code running on more cores
 - **No-one** can debug whole-machine jobs on any of these systems
 - **How many systems next year ...?**




- Scalar features
 - Advanced C++ and STL
 - Fortran 90, 95 and 2003: modules, allocatable data, pointers, derived types
 - Memory debugging
- Multithreading & OpenMP features
 - Step, breakpoint etc. one or all threads
- MPI features
 - Easy to manage groups
 - Control processes by groups
 - Compare data
 - Visualize message queues (not Cray!)



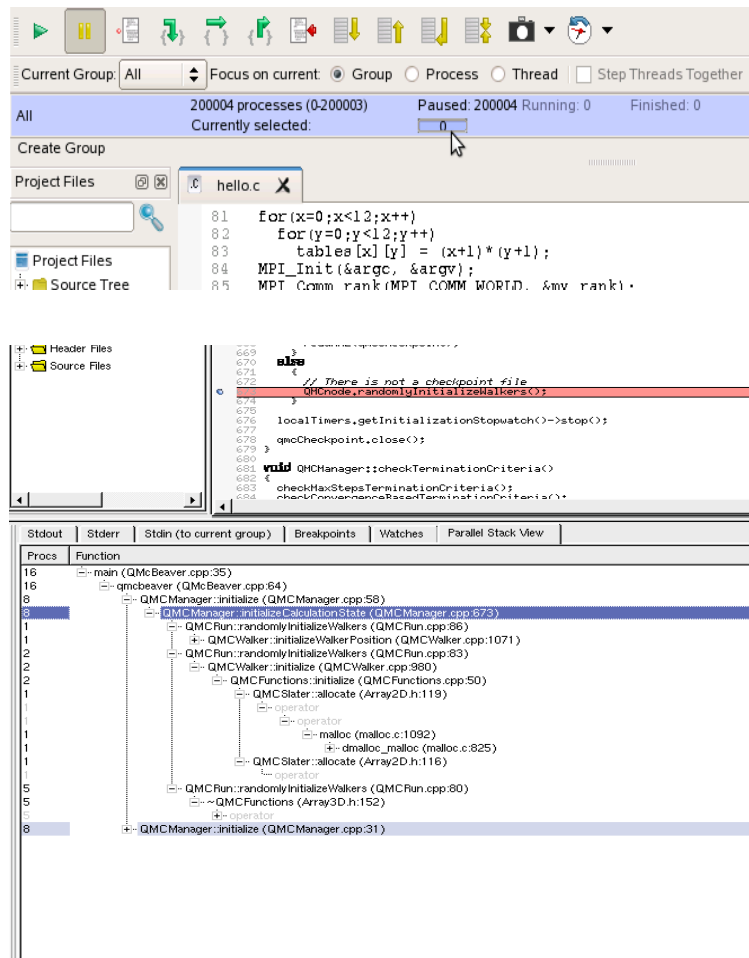
- Many benefits to graphical parallel debuggers
 - Large feature sets for common bugs
 - Richness of user interface and real control of processes
- Historically **all** parallel debuggers hit scale problems
 - Bottleneck at the frontend: Direct GUI → nodes architectures
 - Linear performance in number of processes
 - Human factors limit – mouse fatigue and brain overload
- Are tools ready for the task?
 - Allinea is changing the game!

- Allinea is developing a petascale debugging tool
 - Production Grade debugger
 - Multi-year project – with usable intermediate results
 - Commenced June 2009 – showing results already
- Building a multi-level tree for debugging
 - One tool from 1 to 250,000 cores
 - Goal of logarithmic performance scaling
- Scaling all aspects of debugging
 - Step, attach, data checking, ...
 - Many challenges ahead!

DDT 3.0 Performance Figures



- DDT is delivering petascale debugging **today**
 - A collaboration with ORNL on Jaguar Cray XT
 - Tree architecture – logarithmic performance
 - Many operations now faster at 220,000 than previously at 1,000 cores
 - **~1/10th of a second** to step and gather all stacks at 220,000 cores



- Control Processes by Groups
 - Set breakpoints, step, play, stop etc. using user-defined groups
 - Scalable process groups view
 - Compact representation
- Parallel Stack View
 - Finds rogue processes faster
 - Identifies classes of process behaviour
 - Allows rapid grouping of processes

Expression:

Processes in current group (All, 189120 procs)

☐ Limit comparison to s.f.

☐ Filter:

☒ Align stack frames

| Value | Process(es) | Count |
|-------|-----------------------------------|-------------|
| 0 | 0,3,6,9,12,15,18,21,24,27,30,... | 189120 |
| 1 | 1,4,7,10,13,16,19,22,25,28,31,... | Filtered: 0 |
| 2 | 2,5,8,11,14,17,20,23,26,29,32,... | Errors: 0 |

Aggregate: 0

Numerical: 189120

Sum: 189120

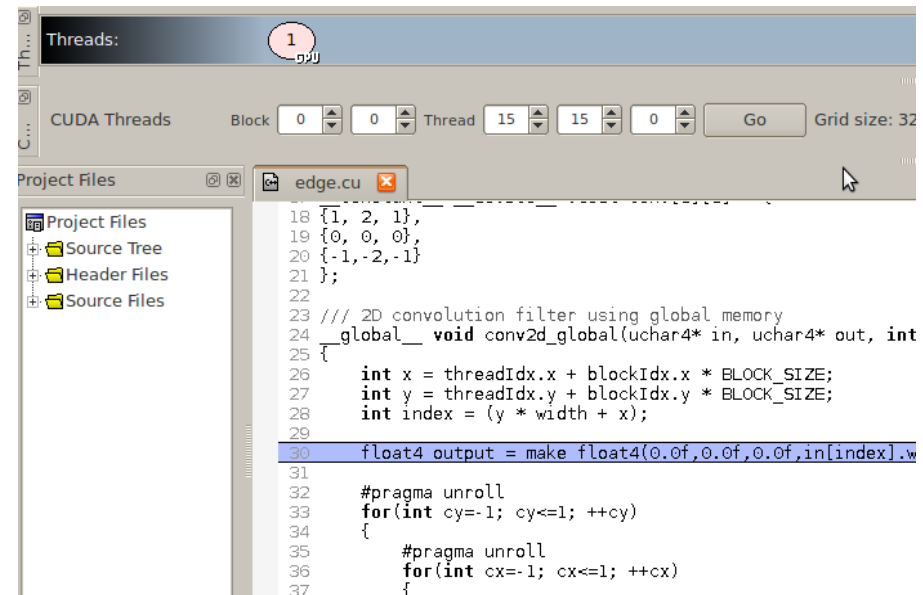
Minimum: 0

Maximum: 2

- Gather from every node
 - Potentially costly – if all data different
 - ... easy if values mostly same
 - New ideas
 - Aggregated statistics
 - Probabilistic algorithms optimize performance – even in pathological case
 - With a fast and scalable infrastructure, new things become possible
 - Watch this space!

- DDT is the first Petascale debugger..
 - A debugging tool has finally caught up with the hardware!
 - Work is in progress to port every feature for scale
 - Memory debugging, data visualization,
 - How can the infrastructure be built upon?
 - Does DDT offer the right framework for collaboration?
 - Can we encourage a codebase of user-generated MPI tools/utilities?
- ... but large clusters are a fraction of HPC
 - Most parallel development starts smaller
 - Is now starting even smaller: GPUs

- Run the code
 - Browse source
 - Set breakpoints
 - Stop at a line of CUDA code
 - Stops once for each scheduled collection of blocks
- Select a CUDA thread
 - Examine variables and shared memory
 - Step a warp
 - View all extant threads in parallel tree view



```

18 {1, 2, 1},
19 {0, 0, 0},
20 {-1,-2,-1}
21 };
22
23 // 2D convolution filter using global memory
24 global__ void conv2d_global(uchar4* in, uchar4* out, int
25 {
26     int x = threadIdx.x + blockIdx.x * BLOCK_SIZE;
27     int y = threadIdx.y + blockIdx.y * BLOCK_SIZE;
28     int index = (y * width + x);
29
30     float4 output = make_float4(0.0f, 0.0f, 0.0f, in[index].w);
31
32     #pragma unroll
33     for(int cy=-1; cy<=1; ++cy)
34     {
35         #pragma unroll
36         for(int cx=-1; cx<=1; ++cx)
37     {
    
```

| Threads | Function |
|---------|----------------------------|
| 1 | main (edge.cu:75) |
| 32 | conv2d_global (edge.cu:35) |
| 480 | conv2d_global (edge.cu:39) |

- DDT 2.4.1 installed and waiting for you!
 - “module load ddt” on Jaguarpf
 - DDT will submit job for you
- Q4 2009
 - Official DDT 2.5 release: some performance improvements
 - ~16-32k cores
 - Private ORNL access to latest DDT development
 - Much faster process control and data comparison – able to reach 100k cores and higher easily
- H1 2010
 - Further development
 - Scalable memory debugging and data export
 - DDT 3.0 stable release with performance to 250,000 cores

allinea